

IN THE CLAIMS

PLEASE AMEND THE CLAIMS AS FOLLOWS:

1. (currently amended) A method of classifying a message transmitted over a network, the method comprising:

maintaining a reputation table in memory, the reputation table including information regarding a plurality of address-domain pairs, each of the plurality of address-domain pairs indicating an IP address and an associated domain of a previously received message, the information regarding each of the plurality of address-domain pairs including a score based on one or more classification variables, the one or more classification variables decaying with time;

receiving the message transmitted over the network and addressed to a recipient; and
executing instructions stored in a non-transitory computer readable storage medium to:

determine an associated domain from which the received message is purported to be sent,

identify that the determined domain appears on a whitelist associated with the recipient,

determine an IP address corresponding to a device from which the received message was relayed,

associate the determined domain with the IP address to create an address-domain pair for the received message;

~~classify~~ assign a score to the received message ~~based on a score assigned to the address-domain pair,~~ the score comprising a ratio of a first classification variable of the address-domain pair to a second classification variable of the address-domain pair as indicated by the reputation table, wherein the score is indicative of spam, and

classify the received message according to whether a common classification appears across a plurality of IP addresses associated with the domain, wherein:

the score is overridden and the received message is classified as good in accordance with the whitelist, the good classification based on a common classification appearing across the plurality of IP addresses associated with the domain, and

~~override the whitelist is overridden and the received message is classified as spam in accordance with based on~~ the score assigned to the address-domain pair, the spam classification based on no common classification appearing across the plurality of IP addresses associated with the domain wherein the received message is classified as spam even though the domain of the received message appears on the whitelist.

2. (cancelled)
3. (previously presented) The method of claim 1, wherein classifying the received message is further based on classification variables associated with another address-domain pair, the other address domain pair having a related IP address or related domain.
4. (previously presented) The method of claim 1, wherein classifying the received message is further based on classifications of other messages associated with the domain of the received message, the other messages further being associated with IP addresses other than the IP address of the received message.
5. (original) The method of claim 1, wherein a plurality of IP addresses is associated with the domain.
6. (original) The method of claim 1, wherein the IP address is associated with a plurality of domains.
7. (original) The method of claim 1, wherein the IP address is a boundary IP address.

8. (original) The method of claim 1, wherein the IP address is preconfigured.
9. (original) The method of claim 1, wherein the IP address is preconfigured to be one hop from a gateway IP address.
10. (original) The method of claim 1, wherein the IP address is learned.
11. (original) The method of claim 1, wherein the IP address is adaptively determined.
12. (cancelled)
13. (previously presented) The method of claim 10, wherein the IP address is a boundary IP address and wherein the boundary IP address is learned by detecting a pattern in a certain number of previously received messages.
14. (previously presented) The method of claim 1, wherein determining the domain from which the received message is purported to be sent includes identifying the stated sender domain associated with the received message.
15. (previously presented) The method of claim 1, wherein the domain is a domain associated with a boundary IP address.
16. (previously presented) The method of claim 1, wherein classifying the received message is further based on consulting a white list.
17. (previously presented) The method of claim 1, wherein classifying the received message is further based on previous classifications made to the address-domain pair.
18. (cancelled)

19. (previously presented) The method of claim 1, wherein assigning the score includes determining a spam ratio.
20. (previously presented) The method of claim I, wherein assigning the score includes determining a spam rate.
21. (previously presented) The method of claim I, wherein assigning the score includes determining an estimated instantaneous spam rate.
22. (cancelled)
23. (previously presented) The method of claim 1, wherein classifying the received message includes giving a classification variable greater weight relative to another classification variable.
24. (previously presented) The method of claim 1, wherein classifying the received message includes giving a classification variable associated with user classification greater weight relative to a classification variable associated with computer classification.
25. (previously presented) The method of claim 1, wherein classifying the received message includes giving an indeterminate classification a fraction of the weight of a good classification.
26. (previously presented) The method of claim 1, wherein the reputation table is indexed by IP address and domain.
27. (previously presented) The method of claim 1, wherein each cell of the reputation table includes information about previous classifications.

28. (previously presented) The method of claim 1, further comprising providing the classification of the received message based on the address-domain pair as input to another classifier.
29. (previously presented) The method of claim 28, wherein the other classifier is a Bayesian classifier.
30. (previously presented) The method of claim 1, wherein classifying the received message is further based on a score assigned to the IP address.
31. (previously presented) The method of claim 1, wherein classifying the received message is further based on a score assigned to the domain.
32. (previously presented) The method of claim 1, further comprising determining that the received message was forged based on the score assigned to the domain.
33. (previously presented) The method of claim 3D, further comprising determining the score assigned to the IP address.
34. (previously presented) The method of claim 31, further comprising determining the score assigned to the domain.

35. (currently amended) A non-transitory computer-readable storage medium having embodied thereon a program, the program being executable by a processor to perform a method for classifying a message transmitted over a network, the method comprising:

determining an associated domain from which a received message is purported to be sent;

identifying that the determined domain appears on a whitelist associated with a recipient of the received message,

determining an IP address from which the received message was relayed;

associating the determined domain with the IP address to create an address-domain pair for the received message;

~~classifying~~ assigning a score to the received message based on a score assigned to the address-domain pair, the score indicative of spam and comprising a ratio of a first classification variable of the address-domain pair to a second classification variable of the address-domain pair, the first classification variable and the second classification variable indicated by a reputation table including information regarding a plurality of address-domain pairs, each of the plurality of address-domain pairs indicating an IP address and an associated domain of a previously received message, the information regarding each of the plurality of address-domain pairs including a score based on one or more classification variables, the one or more classification variables decaying with time, and

classifying the received message according to whether a common classification appears across a plurality of IP addresses associated with the domain, wherein:

the score is overridden and the received message is classified as good in accordance with the whitelist, the good classification based on a common classification appearing across the plurality of IP addresses associated with the domain, and

~~overriding the whitelist~~ is overridden and the received message is classified as spam in accordance with ~~based on the score assigned to the address-domain pair, the spam classification based on no common classification appearing across the plurality of IP addresses associated with the domain wherein the received message is classified as spam even though the domain of the received message appears on the whitelist.~~